Testimony By Mr. Edward W. Felten

Professor, Department of Computer Science

Princeton University

Open the lid of an electronic voting machine and look inside; what you will see is a computer, much like an ordinary desktop PC or Mac. Because they are computers, e-voting machines are susceptible to familiar computer problems such as crashes, bugs, mysterious malfunctions, data tampering, and even computer viruses. The question is not whether we can eliminate these problems - we cannot - but how we will cope with them.

Unlike ordinary desktop computers, e-voting systems are entrusted with the most important process of our democracy - collecting and counting votes - and must perform that process accurately, reliably, accessibly, and securely. Trust in election outcomes is necessary for our electoral system to work, but the political system often does not lend itself easily to trusting relationships. Voting technologies must help to build this trust. Today's e-voting infrastructure is not up to the task, but tomorrow's can be.

Two weeks ago Ariel J. Feldman, J. Alex Halderman, and I released a paper analyzing in detail the security of the Diebold AccuVote-TS, one of the most widely used e-voting systems. The main findings of our study were as follows:

1. Malicious software running on a single voting machine can steal votes with little if any risk of detection. The malicious software can modify all of the records, audit logs, and counters kept by the voting machine, so that even careful forensic examination of these records will find nothing amiss. We have constructed demonstration software that carries out this vote-stealing attack.

2. Anyone who has physical access to a voting machine, or to a memory card that will later be inserted into a machine, can install said malicious software using a simple method that takes as little as one minute. In practice, poll workers and others often have unsupervised access to the machines.

3. AccuVote-TS machines are susceptible to voting-machine viruses - computer viruses that can spread malicious software automatically and invisibly from machine to machine during normal pre- and post-election activity. We have constructed a demonstration virus that spreads in this way, installing our demonstration vote-stealing program on every machine it infects.

4. While some of these problems can be eliminated by improving Diebold's software, others cannot be remedied without replacing the machines' hardware. Changes to election procedures would also be required to ensure security.

Our web site at http://itpolicy.princeton.edu/voting has links to our full technical report and a ten-minute video showing our demonstration vote-stealing virus in operation. The technical report goes into considerable detail and includes a discussion of why existing election procedures are not sufficient to prevent virus attacks. While we are not alleging fraud in any specific past election, our results do raise serious concern about the security of future elections.

One lesson of our study is that security depends on getting the technical details right. A security measure that sounds robust in the abstract may be useless or worse if implemented poorly. Too often, the designers of the AccuVote-TS failed to get the details right.

A good example is the AccuVote-TS access door. The access door on this machine protects the removable memory card that stores the votes, so the door should be locked securely and access to the keys should be strictly limited. In fact, the tens of thousands of AccuVote-TS machines can all be opened with the same key, and this very same key is used widely in office furniture, jukeboxes, and even hotel minibars. I bought several keys on the Internet from an office furniture shop and a jukebox supply shop, and they all open the AccuVote-TS. Details matter. It is not enough to have a key; it matters which key you use.

Some voting machines, including the AccuVote-TS, record votes internally in a computer file, with the votes stored in the order they were cast. This approach endangers the secrecy of the ballot. If election procedures record the order in which voters cast their votes (or allow partisan observers to do so, as is the practice in my polling place), then a sequential record of the votes can be correlated with the order of voters to reconstruct the ballots cast by individual voters. The AccuVote-TS is one voting machine that gets this detail wrong.

The AccuVote-TS suffers from many such problems. It encrypts stored votes, but stores the secret decryption key where it is easily found by hostile software. It keeps two redundant copies of each stored vote, but both copies are subject to easy tampering. Some of these errors are

more technical in nature than the access-door key error and the vote-recording error, but they are just as serious.

The implications of our study go beyond the specific voting machine we studied to reveal broader systemic problems. More worrisome than any specific vulnerability is that, despite its many problems, the system we studied was certified, purchased and deployed by many states and counties, and is slated for use in the upcoming November election. This leads us to conclude that existing certification and procurement procedures are inadequate to prevent the kinds of serious vulnerabilities we discovered. Here again the details matter, and too often current processes get the details wrong.

Though some claim that election procedures will prevent the kinds of problems we identified, the rigid procedures described in vendor manuals are often ignored in practice. Machines are supposed to be sealed with numbered security tape; but missing or broken tape is usually ignored, and election workers often break the tape themselves when trying to revive malfunctioning machines. Machines and removable vote-storage media are theoretically kept under lock and key, but in practice they are often sent home with election workers or left unattended. At my polling place in Princeton, the night before an election, the DRE machines sit unattended in an unlocked elementary school lobby where anyone could tamper with them. Stringent official procedures only matter if they are followed in practice.

There are several things we can do to improve the security of our e-voting infrastructure.

In the short term, some limited steps are still feasible before November. Given the susceptibility of some e-voting systems to electronic tampering, we should take extra care to secure the chain of custody for voting machines and vote-storage media from now until Election Day. This cannot repair machines that have already been tampered with, but it can reduce the likelihood of further tampering. Needless to say, what we need is not more memos laying down theoretical procedures, but detailed execution to narrow the gap between procedural theory and practice.

In the medium term, I offer three recommendations. First, we should fix the certification process to better account for security. Certification seems to focus on machine attributes that are easily tested, but

security problems are difficult to detect by testing because no predetermined set of test scenarios can account for the tactics of a clever adversary who systematically exploits gaps in a system.

In practice, the certification process often misses security problems that are simple but very dangerous. For example, the AccuVote-TS system we studied will silently accept and install any software update offered by any memory card that is inserted into the system. The system makes no effort to verify that the offered update is authorized by the vendor, election officials, or anyone else. This is a very serious weakness that opens the door to the injection of malicious software and the silent, automatic spread of viruses. Yet the system was certified despite this obvious vulnerability. The existing certification process seems unable to detect such problems reliably. It must be improved.

Second, a voter-verified paper audit trail (VVPAT) is a necessary safeguard given the state of the art today. With these paper trails, as with other voting technologies, we must get the details right - poorly designed paper trails can be unreliable or hard to use, or can compromise the secrecy of the ballot - but a well-designed paper trail can improve security and enhance voter confidence, without compromising accessibility.

In comparing VVPATs with paperless DREs, we must compare apples to apples. For example, we must not compare a VVPAT that compromises the secret ballot by recording votes in the order cast (e.g., on a continuous roll of paper) with a paperless DRE that gets this detail right. Instead, we must assume good engineering in both cases, and weigh the significant security benefits of VVPATs against their costs.

Paper records, either VVPATs or traditional paper ballots, have their drawbacks. They are not immune to fraud. What is important is that they have different failure modes than electronic records, so that the combination of electronic and paper recordkeeping, if implemented well, can be more robust against fraud than either would be alone.

One aspect of a well-implemented VVPAT system is that the electronic and paper records must be compared to each other. We do not need to verify every paper record, just enough to detect large-scale fraud. Unless an election is very close - which will probably trigger a full recount anyway - checking a few percent of ballots will suffice. Similarly, it is not necessary for every voter to read and verify the paper record of his vote; as long as even a few voters do so, any tampering widespread enough to be significant will be easily detected.

Third, we must do more to leverage the expertise of independent security experts. Independent analyses, by experts neither paid by nor reporting

to voting machine vendors, have discovered many areas for improvement in today's technologies, yet most vendors systematically try to prevent such analyses. For example, my colleagues and I would be happy to examine other versions of Diebold's AccuVote-TS or AccuVote-TSx software to determine whether they are subject to the vote-stealing virus problems we have identified; but Diebold refuses to let election officials call on us for this purpose. Other vendors follow a similar policy of resisting public study and discussion of the technologies that count our votes.

In the long run, further research is needed to help us understand how to improve the voting system. For example, fully electronic verification technologies may one day be a viable substitute for VVPATs, once researchers have worked out the details necessary to deploy them in the real world accessibly and securely. We also need more systematic studies of what really happens in polling places, especially when problems arise. Finally, there is much to learn from work in other areas of computer security - today, even video game consoles like the Xbox are more tamper-resistant than voting machines.

Those not versed in computer security can miss the significance of e-voting security vulnerabilities. From a security standpoint, what distinguishes computerized voting systems from traditional systems is not that computers are easier to compromise, but that the consequences of compromise can be so much more severe. Breaking into an old-fashioned ballot box can affect a few hundred ballots at most; injecting a virus into a single computerized voting machine can affect an entire election.

Intuitions developed with older technologies can mislead when applied to computerized systems. For example, non-experts often fail to appreciate how difficult it is to tell what is happening inside a computer system. We cannot "just look" to see what is happening or whether the right software is installed. Often our only recourse is to ask the system itself what it is doing - which is fine if the system is working correctly, but fruitless if the system is compromised. There is no point in asking a virus whether a virus is present.

Similarly, non-experts often assume that pre-election testing is an effective way to trigger and detect malicious software that might have infected a voting machine. Here again, computerized systems are different. A modified lever machine will work the same whether or not it is Election Day; but malicious software on a DRE can check whether the machine is in pre-election testing mode, or can check the date, or can check whether the number and pattern of voters is consistent with election day, and can activate its vote-stealing capability only in a real election. Our demonstration AccuVote-TS virus takes measures to remain inactive and thus evade detection during pre-election logic and accuracy testing. It is very difficult to tell whether such a virus is present. In general, malicious software is much harder to detect than non-experts would expect.

My point is not that these challenges are insurmountable but that one needs specialized knowledge and sophisticated analysis to figure out what is possible. Acknowledging that security experts can learn from election experts, I submit that election experts can also learn from security experts.

Getting the details of voting right is difficult, especially in today's high-tech polling place. But failure is not an option. The stakes are too high, and the risk of malfunction or fraud too great, to make our current course tenable in the long run. We need to work harder and smarter, exploiting the knowledge of both election experts and technical experts.

Biography of Edward W. Felten

Edward W. Felten is Professor of Computer Science and Public Affairs, and Director of the Center for Information Technology Policy, at Princeton University. His research interests include computer security and privacy, Internet software, and information technology policy. He has published more than eighty papers in the research literature, and two books, and he is widely quoted in the press as an expert on security, privacy, and information technology policy. He has advised the U.S. Departments of Justice, Defense, and Homeland Security, and the Federal Trade Commission, on security-related issues. He serves on the Executive Committee of USACM, the U.S.
public policy committee of ACM, the leading professional society for computer scientists. In 2003, Scientific American magazine named him to its list of fifty global leaders in science and technology.